Generative AI in AgriFood

By A M Howcroft, with assistance from AVA

March 2023

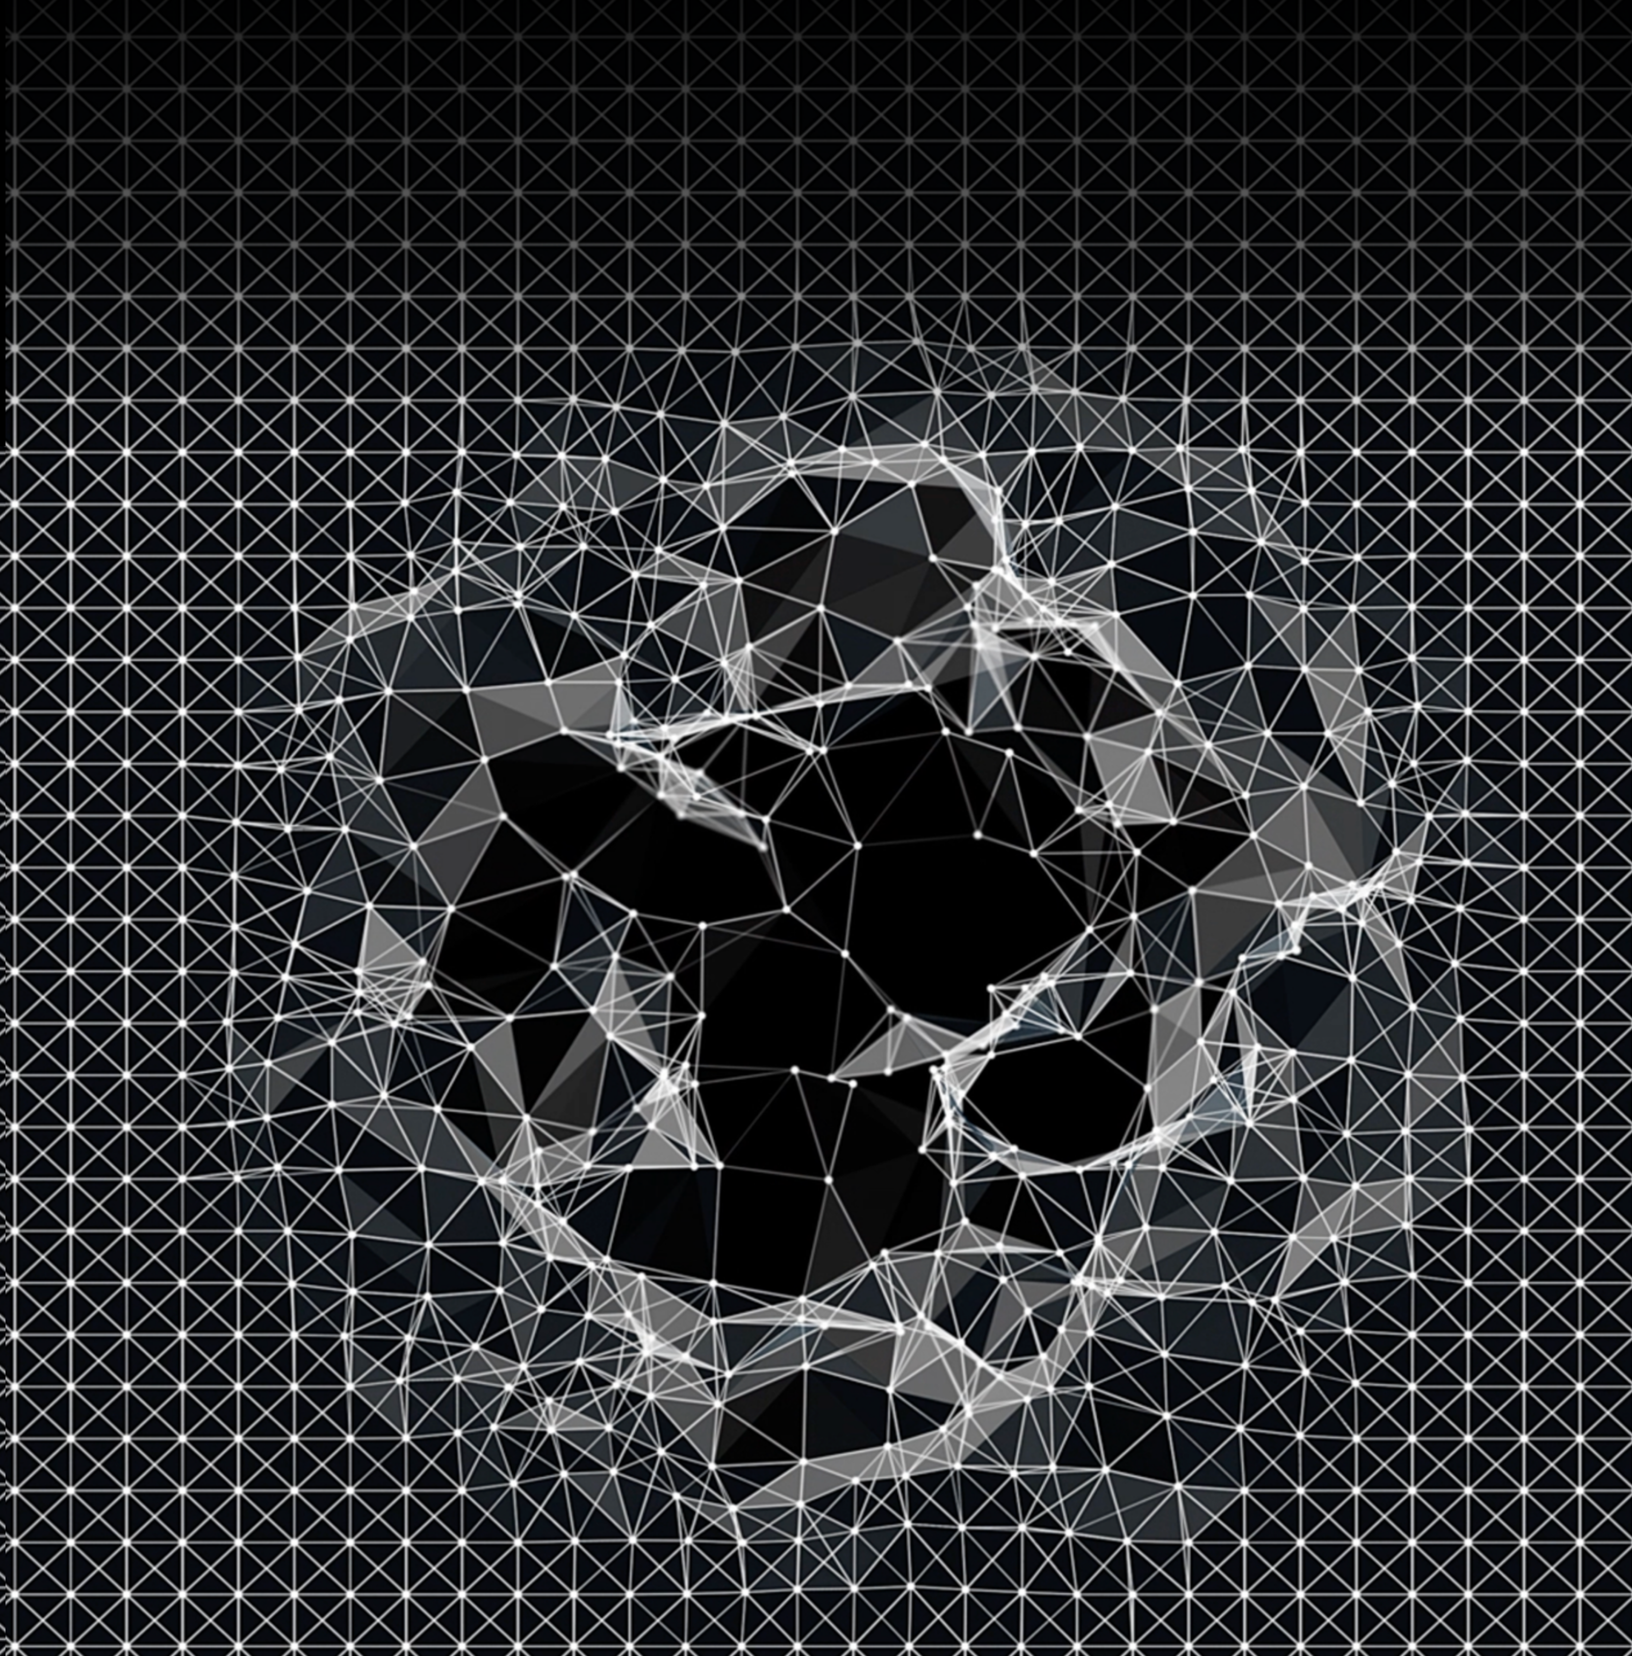

The future is already here,
it’s just not evenly distributed

William Gibson

OVERVIEW

Exec Summary

Generative AI, particularly Large Language Models (LLMs) like ChatGPT, are a hot topic in the media, and being deployed by early adopters. The global generative AI market is predicted to reach $109.37 billion by 2030, with a CAGR of 34.6%. LLMs will have a significant impact on multiple functional areas and industries, and AgriFood will be no exception. LLMs use machine learning algorithms to generate human-like responses by analyzing massive amounts of data. However, the rapid progress of generative AI raises ethical and legal concerns due to potential bias, and the role of human judgment in decision-making remains crucial.

LLMs have ten functional capabilities that apply across different industries, including natural language understanding and generation, translation, text summarization, and text simplification. The most common early implementations are in functions such as customer support and service, marketing and advertising, sales and business development, legal, and human resources. In addition to these functional applications, there are specific AgriFood scenarios where the LLM approach can be usefully deployed, such as precision agriculture, supply chain optimization, food safety, and sustainability. We also review the primary limitations of the technology, such as the alignment problem and issues with data freshness, data bias, unintended consequences, interpretability, and the intensive resources required.

We introduce AVA, the AgriFood Virtual Advisor by SWARM Engineering, as an example of the next wave of specialized agents with industry training that enhances the interaction between a vendor, an organization's technology platform, and the end-user.

The report discusses the importance of questioning skills in the new era of LLMs, as AI agents answer queries and generate content based on prompts. The quality of their responses will depend largely on a user’s expertise at questioning. We suggest improving organizational questioning capabilities through training and executive support.

We predict a backlash against LLMs due to concerns about privacy, bias, ethical considerations, job loss, and environmental impact. AgriFood companies should address these concerns proactively, using explainable AI with human oversight. We also describe the coming proliferation of AI agents and discuss alternative approaches to manage AI employees.

Successful AgriFood companies of the future will be those that collect and manage high-quality data, incorporate LLMs, innovate, and differentiate themselves through AI. Companies that adopt a ‘sit and wait’ approach will place their business at risk. We provide a checklist to help organizations measure and plan their approach to deploying generative AI and LLMs.

Introduction

Generative Artificial Intelligence (AI) and in particular Large Language Models (LLMs) have become one of the most widely discussed, reviewed, and adopted technologies (at least in experimental form) since ChatGPT was made publicly available in November 2022. They offer the potential to transform the way we interact with technology and perform business. LLMs are artificial intelligence systems that can analyze vast amounts of data, recognize patterns, and generate responses that mimic human language. They are already being used across various industries, from finance to healthcare, to drive efficiency, improve accuracy, and enhance user experiences. The AgriFood industry is no exception, and LLMs will play an increasingly important role in addressing some of the most pressing challenges faced by this industry.

AgriFood, which encompasses agriculture and the entire food supply chain, is vital to global food security and human health. However, it faces many obstacles, including resource scarcity, climate change, biodiversity loss, geo-political risk, and food safety issues. In recent years, technology has emerged as a key tool to tackle these challenges, and generative AI will soon be at the forefront of this technological revolution. By leveraging the power of natural language processing and machine learning, this latest wave of AI has the power to dramatically change AgriFood in numerous ways. However, this rapid progress also raises ethical and legal concerns, and there have been several well documented instances of potential bias in LLMs. The role of human expertise and judgment in the decision-making process cannot be overlooked. This report will explore the impact of generative AI, and specifically LLMs on the AgriFood industry, highlighting the benefits and challenges of this technology and examining the potential future developments and their implications.

Overall, the global generative AI market size is anticipated to reach USD 109.37 billion by 2030, according to a report by Grand View Research, Inc. The market is expected to expand at a CAGR of 34.6% from 2022 to 2030. There will be some very specific industry applications, and many broader uses that span multiple horizontal functions. For example, Allied Market Research reported that in January 2023, Nvidia, in partnership with Evozyne, a pharmaceutical startup, launched a new generative AI model capable of producing proteins for use in medicine and other industries. This new protein transformer variational auto-encoder (ProT-VAE) is built on Nvidia's BioNeMo framework and uses generative AI to rapidly create synthetic protein designs that fit into the given parameters. We expect to see similar niche applications emerge in AgriFood, but the short-term opportunity, estimated at >41% of the market, is for the deployment of LLMs across multiple functional areas.

The pace of change is also accelerating, with ChatGPT-4 being launched in March, showing improved test scores and the ability to handle visual inputs. Google launched its own Bard LaMDA (Language Model for Dialogue Applications) service shortly afterwards, differentiating with live internet data feeds, and multiple answers or 'drafts' to questions. We expect to see rapid iterations of the primary LLMs as they fight for market position.

Forms Of Generative AI

Generative AI refers to artificial intelligence systems that can create new and original content, such as text, images, or videos, rather than simply making predictions or classifications based on existing data. These systems use machine learning algorithms that learn to generate new content by training on large datasets of existing material. There are various types of generative AI systems, such as:

GPT
Generative Pre-trained Transformer

A neural network architecture developed by OpenAI that can be fine-tuned to generate text, images, or other content

VAE
Variational Autoencoder

Another neural network architecture that can be used to generate new data based on patterns learned from existing data

GAN
Generative Adversarial Network

An architecture that generates new data by pitting two neural networks against each other in a game-like setting

RNN
Recurrent Neural Network

Another style of neural network that can generate sequences of data, such as text or music, based on patterns learned from existing sequences

In recent months numerous generative AI tools have appeared to cover a range of tasks such as generating artwork in the style of a specific artist, creating logos, writing business plans, creating presentations, or even having an AI generated character do a photoshoot.

There have been some darker examples such as the use of generative AI in the creation of deepfakes, which are false media that look or sound like the real thing. Deepfakes can produce highly convincing images, videos, or audio recordings of people or events that never happened. While they can be used for harmless fun, such as creating entertaining videos or memes, they can also be malicious, with nefarious purposes such as spreading misinformation, propaganda, non-consensual pornography, or fake news. The rise of deepfakes has led to concerns about their impact on public trust in media and the potential for them to manipulate opinions. As a result, there have been efforts to develop technologies to detect deepfakes and prevent their dissemination. This is the age-old problem of technology being capable of delivering both positive and negative results, depending on the person wielding the tool.

While there is a broad range of generative AI capabilities, the rise of LLMs with GPT, such as ChatGPT from OpenAI and Bard from Google are at the forefront of driving change for most business organizations and will be our focal point moving forward.

If you are interested in learning more about other forms of generative AI, we would encourage you to look at tools such as Jasper AI (generates articles and copywriting pieces), DALL-E (the image generator from OpenAI) Midjourney or Craiyon as an alternative for images, Dadabots or AIVA for music generation, and tools like GitHub co-pilot or Codestarter if you want to try your hand at generating software.

> "As more and more artificial intelligence is entering into the world, more and more emotional intelligence must enter into leadership."
>
> *Amit Ray*

CAPABILITIES & LIMITATIONS

Technical Capabilities

Before we explore the functional capabilities of LLMs, let's take a moment to consider the technical aspects. For the lay person, an LLM can give the appearance of intelligence, but the systems are not actually capable of thinking. In the words of Sam Altman (CEO of OpenAI) *"ChatGPT is incredibly limited, but good enough at some things to create a misleading impression of greatness"*. The primary technical capabilities of LLMs include natural language processing, deep learning, and language modeling, which combined make them powerful tools for analyzing and processing large amounts of unstructured data (such as books, emails, or academic papers), generating insights and predictions, and communicating with humans in a way that feels natural and intuitive. Here's a brief overview of their technical capabilities:

Natural Language Processing (NLP)

A field of AI that focuses on the interaction between computers and human languages, NLP allows LLMs to understand and analyze human language, including grammar, syntax, and semantics to identify patterns, extract insights, and generate responses that mimic human language

Language Modeling

Predicting the probability of the next word in a given sequence of words. Language modeling allows LLMs to generate text that is coherent and contextually relevant. By using language modeling, LLMs can understand and respond to human language in a seemingly natural way

Deep Learning

A type of machine learning that uses artificial neural networks to analyze and interpret complex data. Deep learning allows LLMs to learn from large amounts of data and improve their accuracy over time. By using deep learning, LLMs can recognize patterns, classify information, and make predictions based on historical data

In addition to these primary capabilities, LLMs also have several other important features that enable them to perform complex tasks. These features include:

- **Attention Mechanisms:** which allows them to focus on specific parts of a sentence or document when processing information.
- **Transfer Learning:** the ability of LLMs to transfer knowledge learned from one task to another, which can improve their efficiency and accuracy.
- **Encoder-Decoder Architecture:** a type of architecture that allows LLMs to process inputs of varying lengths and generate outputs of varying lengths, making them highly flexible and adaptable.

Training

We will look at ChatGPT as a specific example of how an LLM is trained, since there is a wealth of public data available on this topic. ChatGPT was trained using a combination of unsupervised and supervised learning techniques, with 570GB of text data – or the equivalent of several hundred thousand books. The training data actually consisted of a diverse set of internet text, including web pages, books, and social media posts.

The number of parameters in the GPT model has increased significantly with each version, with ChatGPT-3 having 175 billion parameters, making it the largest language model at the time of its release, and ChatGPT-4 is believed to have 1 trillion parameters. This allows ChatGPT to capture more nuanced relationships between words and generate more contextually appropriate responses. For clarity, we should explain that in machine learning a parameter refers to a variable learned by the model during the training process. A parameter typically refers to a single numerical value in the model's weight matrix. These weight matrices are used to represent the relationships between different elements in the data, such as the relationships between words in a sentence or between topics in a body of text. If ChatGPT were evaluating the sentence "The cow jumped over the", it might assign a high weight value to the word "moon" based on the patterns it has learned, to influence the model's prediction for the next word in the sentence. Other words with a high weight factor could include "fence", "gate", "barn", or "stream", and the choice would depend on the context, and also the training data – if there were lots of nursery rhymes, then moon might be likely, while if the training data was more skewed towards veterinary reports, we might find another word taking precedence.

In AI terms, ChatGPT is a dense LLM (as opposed to a sparse model). This means that it is more expensive to train but is more accurate as it uses a deeper neural net to process each request. Alternative sparse models have a shallow neural net with less parameters but have the advantage of being faster to run (and can potentially run on mobile devices), cheaper to train, and the AI is simpler to understand or 'explain' if there is an issue with the output. Both LLM model types are likely to find their own niche in different applications.

In addition to unsupervised learning, where the system learns by itself, ChatGPT was also fine-tuned using supervised learning (i.e. with human input) on specific tasks, such as language translation, question-answering, and sentiment analysis. Fine-tuning involves adapting the pre-trained model to a specific task by training it on a smaller dataset of labeled examples. This fine-tuning process allows ChatGPT to perform specific tasks more accurately and efficiently.

While models like ChatGPT come 'pre-trained' an organization has the ability to take the core model and train it with specific additional inputs. For example, you could train the LLM to understand pollination processes and results on apple trees, using a combination of historical data, academic papers, and more. The LLM could then answer very specific questions about this functional area. Different pricing models exist for training (expensive) vs using a dataset (cheaper). Training allows an organization to adapt an LLM to their specific requirements.

Limitations

For all of the powerful tasks that LLMs can perform, they do have several limitations that can impact their performance and accuracy. It is good to be aware of these, and to understand the implications.

One of the most pressing issues is commonly referred to in the AI community as **the alignment problem**. This refers to the challenge of ensuring that the goals and values of LLMs align with the goals and values of the humans who create and use them. This problem arises because LLM systems are designed to optimize specific objectives, such as accuracy in natural language processing, and do not comprehend the content, i.e. they are not actually 'thinking' machines. They generate content without regard for broader ethical or social considerations. This is a complex challenge that requires a multi-disciplinary approach with researchers, policymakers, and industry stakeholders. Addressing this problem may involve developing new training methods, improving transparency and interpretability, and establishing guidelines and standards for the ethical use of AI. It is not a problem unique to generative AI, and in fact the significance of this challenge will become far more urgent when we develop 'general AI' systems with the potential for super-intelligence, as we shall see later.

Another major challenge is the **freshness of the data**. Given that LLMs can be expensive to train, the systems may only be trained intermittently, resulting in out-of-date information. ChatGPT-3 famously has data only until the end of 2021, for example, although Google Bard can pull live data from the Internet. There is a related issue, which results from older data being used in training. If you were to ask for a medical diagnosis, and the system had a training dataset of ancient academic papers, you may be offered an approach that current medical practice would not recommend – the LLM, remember is looking at how frequently words occur together – it does not comprehend the answer. USA today quoted an example where someone could ask an LLM, "what's the best treatment for diabetes?" and a likely response would be the diabetes drug "metformin" – not because it's necessarily the best but because it's a word that often appears alongside "diabetes treatment". This is where fine-tuning an LLM can be critical. Other than these issues, here are the most common limitations:

Data Bias	as they are trained on large datasets of text data, which may contain biases that reflect the societal and cultural views of the individuals who created the data, LLMs can perpetuate these biases in their outputs. This can have harmful consequences, especially for marginalized groups, resulting in damaging publicity for all involved
Unintended Consequences	LLMs are designed to optimize specific objectives such as accuracy in natural language processing, without regard for broader ethical considerations. This leaves them open to manipulation, to create polished text to spread false information, and/or manipulate people
Interpretability	as complex neural networks with many layers and parameters, it can be very difficult for humans to understand how LLMs arrive at their outputs. This can be problematic when they are used in high-stakes applications where the outputs of the model can have significant consequences
Resource Intensive	LLMs require a lot of computing power and memory to train and run, which can be expensive and time-consuming. This can be a barrier for smaller organizations or researchers without access to high-performance computing infrastructure
Limited Context	LLMs are trained on specific datasets, which can limit their understanding of any broader context. This can lead to errors or inaccurate outputs when presented with new information outside of their training data

Many writers have rejected LLM tools as ‘toys’ that are not ready for the mainstream. They frequently list the issues described above; the systems can be biased, wrong, decisions can be hard to interpret, and so on. I would not dispute the points, but I would counter that argument by asking the journalists to reflect on their dealings with real people. I have found they can often be wrong, biased, they may not have the full context, and in many cases I can’t fathom their decision process. How many times have you seen or heard a story and thought ‘What were they thinking?!’ It seems to me that LLMs do not have an exclusive position as flawed actors in our lives, but they do have many strengths and advantages, despite their limitations.

Functional Capabilities

While generative AI has a broader range of capabilities, such as image generation or manipulation, an LLM typically has the following ten capabilities:

1. **Natural Language Understanding (NLU):** to understand natural language input (either typed or spoken) and perform tasks such as text classification, sentiment analysis, and named entity recognition.
2. **Natural Language Generation (NLG):** to generate natural language output, such as text, speech, or dialogue, in response to a given input.
3. **Language Translation:** to translate from one language to another.
4. **Text Summarization:** summarize longer texts, allowing for easier consumption and understanding of complex information, e.g. a medical record
5. **Text Simplification:** simplify complex or technical language, making it more accessible to a broader audience.
6. **Answer Questions:** respond to questions posed in natural language, either by providing a direct answer or by extracting information from other data sources.
7. **Dialogue Generation:** by remembering context, an LLM can generate a convincing dialogue, allowing for more engaging and interactive conversations.
8. **Text Classification:** can be used to classify text into different categories or topics, for applications such as spam filtering, or intent detection.
9. **Sentiment Analysis:** analyze the sentiment or tone of a given piece of text, allowing for insights into social media opinions or customer feedback.
10. **Specializations:** trained expertise in certain areas, e.g. ChatGPT can actually write code in specific coding languages

Applications For Business

LLMs are already making a significant impact on various industries, revolutionizing the way businesses interact with customers, automate processes, and make decisions. The functional capabilities apply across different industries, from finance through to healthcare, retail, and e-commerce, to help drive innovation and growth.

In finance, for example, LLMs are being used for fraud detection, risk assessment, and personalized financial advice. While in healthcare, LLMs are being cautiously deployed for medical diagnosis, drug discovery, and patient engagement.

Microsoft has already deployed this technology in GitHub, the code management platform, and announced that ChatGPT functionality will be added to its search engine Bing and is expected to be available in new iterations of Microsoft Office later this year. Salesforce and HubSpot are two of the sales & marketing platforms to have already delivered OpenAI integration, and as the technology continues to evolve, we can expect to see even more significant impacts on industry and society.

Typical cross-industry functions where LLMs can be deployed for high impact, that apply equally in AgriFood, include:

- **Customer support and service:** automate customer service inquiries and provide personalized recommendations and support.
- **Marketing and advertising:** generate copy, customer insights, predict customer behavior, and optimize advertising campaigns
- **Sales and business development:** automate lead generation, prioritize sales opportunities, and provide sales representatives with real-time insights.
- **Supply chain and logistics:** support efforts to optimize supply chain operations, predict demand, and reduce waste.
- **Human resources:** automate HR tasks, such as resume screening and candidate matching, and provide personalized employee feedback and development.
- **Finance and accounting:** fraud detection, risk assessment, and financial forecasting.
- **Legal:** contract analysis, legal research, and document review.
- **Education and training:** for personalized learning, educational content creation, and student engagement.
- **Social media and content creation:** generate content, predict audience engagement, and optimize and review social media campaigns.
- **Management reporting:** automated production of summary reports, natural language query of business information, identification of challenges

These may be good areas to look at first within your organization, as there are many successful examples to learn from.

AgriFood Scenarios

In addition to the standard functional application discussed in the previous section, there are also specific AgriFood scenarios where the LLM approach can be deployed to great effect – here's a limited selection to highlight the possibilities. Clearly, the opportunities will vary based on whether you are a grower, inputs provider, retailer, or food producer.

Crop Monitoring & Management

To analyze sensor data to monitor crop health, predict yields, and optimize fertilizer and irrigation schedules

Pest & Disease Detection

Detect the presence of pests and diseases in crops by reviewing unstructured and structured data sources so growers can respond quickly to outbreaks

Precision Agriculture

Analyze data from sensors and other sources to help optimize planting, irrigation, and fertilizer application

Animal Health Monitoring

Analyze data to monitor the health and behavior of livestock. This can help farmers detect and respond to health issues more quickly

Market Intelligence

Use market and social data to provide insights into consumer trends and preferences, to help growers and food producers make better decisions on planting and selling

Food Safety

Monitor food safety and quality data throughout the supply chain, reducing the risk of foodborne illness and improving consumer confidence

Soil Analysis

Review soil sample reports and provide recommendations for soil management, helping farmers optimize soil health and fertility

Sustainability Metrics

Analyze data from disparate systems to calculate ESG metrics such as carbon footprint, water usage, and biodiversity impact to move towards regenerative ag

These are a selection – the reality is that almost every AgriFood process has potential to benefit.

AI Employees

The Proliferation of AI Agents

As generative AI and LLMs become embedded in many additional functions and roles within both our organizations and daily lives, there will be many ways in which we interact with them. Some of these will be virtually indistinguishable from today, with improved quality. For example, the suggestion of a different layout or additional content in a PowerPoint slide, or a more eloquent and valuable response to a Bing or Google search.

There are also situations where we are already conditioned to interacting with a chatbot, and these (should) simply get better: the dialogue with a text-based support agent at your bank, where they actually resolve the issue instead of handing you off to a human, or a conversation with Siri or Alexa where they actually answer the question instead of suggesting you 'look at this webpage'.

Several applications will add buttons to automatically generate content with or without prompts for your preferences – look for this in email software, CRM systems, social media platforms and even in technical areas such as Microsoft Visual Studio where it will potentially offer increasingly powerful code snippets.

We predict there will be a big rise in the number of virtual agents, avatars, and digital humans in the market. These will combine video, audio, and text to offer specialized services. The latest developments in gaming technology, voice recognition, and real-time translation of text-to-speech will result in digital agents that look like high-quality CGI characters from a Hollywood movie, who can interact with you by using your preferred method of communication, either via spoken or written word. The last few years has seen a proliferation of these technologies, from firms such as Synthesia, Colossyan, Uneeq, Murf.ai, and many more. These digital agents can be the visual interface to an underlying LLM, and are already capable of high-quality dialogue, in multiple languages. The capabilities of LLM will mean that these technologies will suddenly accelerate in the market, especially for firms that can deliver real-time text to video animation for a realistic digital human.

With the correct credentials at your organization, the agents will be able to perform tasks such as:

- Communicate with employees
- Book meetings in your calendar
- Send emails
- Review / summarize / translate your emails
- Create content (writing emails, reports, etc.)
- Keep you informed of critical updates
- Discuss corporate data
- Generate Business Intelligence reports
- Take notes at meetings
- Participate in voice and video calls
- Give you summaries of customer trends

In fact, many of you may already be using agents to perform one or more of these tasks today. Expect the capability, reach, and ease of interaction for each of these agents to take a significant step forward in the next year, or at least be prepared to replace them with competing firms that offer such alternatives.

> “The only way you can predict the future is to build it”
>
> *Alan Kay*

AVA – the AgriFood Virtual Advisor

In addition to the standard capabilities that will come baked-in for any LLM trained agent, each digital assistant may come with specialized training and capabilities. As an example of this, at SWARM Engineering we announced *AVA – the AgriFood Virtual Advisor*. She will offer the core capabilities listed in the previous section, and will also be equipped to:

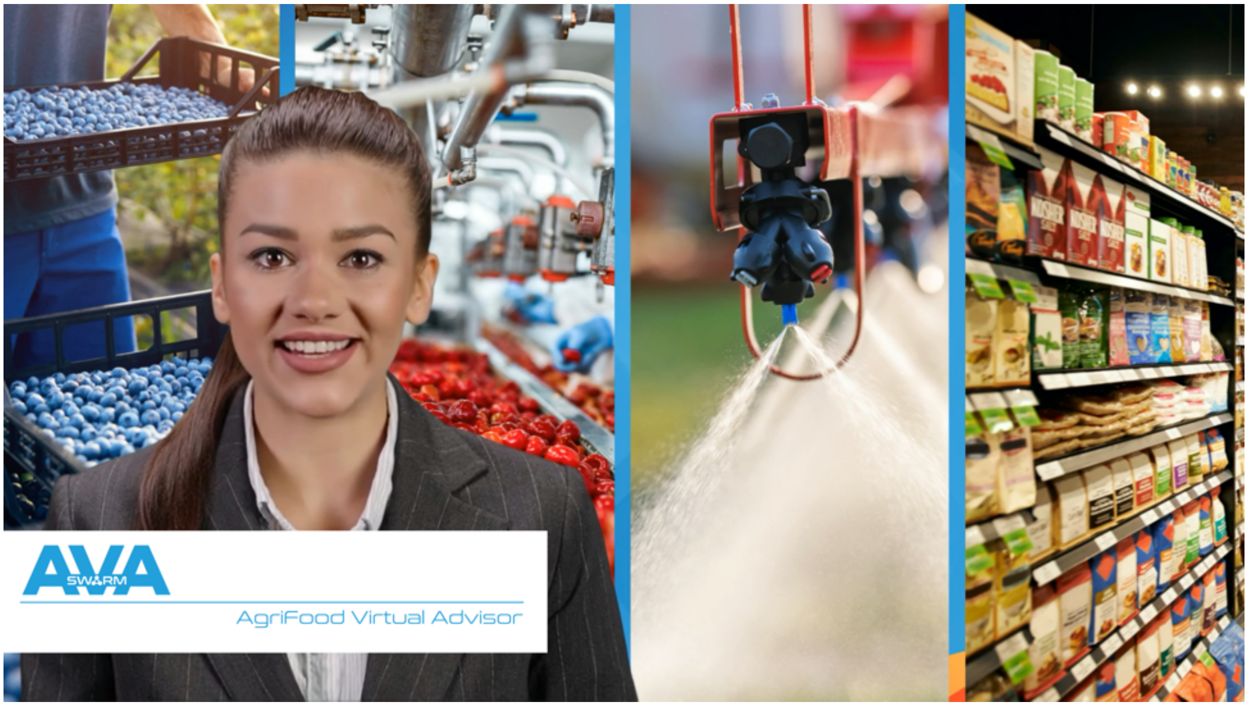

1. **Understand AgriFood Terminology**
2. **Identify, define, and solve operational challenges**

In practice, what that means is that AVA is trained with an additional lexicon of agricultural jargon, for ease of communication with our customers, and she can interact directly with the underlying SWARM Challenge Management and Solution Engine technology. Essentially, AVA enhances the interaction between our underlying technology platform and the end user. We see this as a major step forward in productivity, as it makes it much simpler for an end user to describe their problem, in a conversational format, which we directly translate into a full SaaS solution running advanced AI or Operational Research algorithms.

We have been able to achieve this because we already possess the required training materials for a private instance of ChatGPT. Namely, a library of templates and worked examples that cover traditional challenges in the AgriFood industry, such as balancing supply and demand, labor planning, optimizing inbound and outbound logistics, inventory management, network optimization, and many more use cases (we have 40+ within the platform, growing every month). By training AVA on the structure and content of our templates, and the related terminology, with her ChatGPT LLM she can hold conversations - in parallel, we might add - with the many people involved in an operational process. She can discover the scale of the challenge, how it is currently solved, and the business impact a solution could deliver. AVA can work with the team to identify process goals and metrics, along with the levers that can be used to affect the outcome. She captures critical constraints that often derail a project, plus the hidden tribal knowledge that can be so hard to gather, and automatically turns this into a Challenge Definition report which can be shared with the team, and the business and technical sponsors.

Our experience shows that having a standard structure to define a challenge is helpful to both people and machines, resulting in clear communication, and a process that continually improves. We actually take the underlying meta-data from the document and match this directly to a solution. Which means that after AVA has worked with an organization to define a challenge, we can rapidly deliver a fully working solution.

It's a powerful combination, and ultimately we see no reason why AVA couldn't be trained to capture requirements and deliver solutions for our partners, too. In fact, the more knowledge and training we put into AVA, the more valuable she becomes to our customers, our partners, and ourselves. It is a classic case of a virtuous circle.

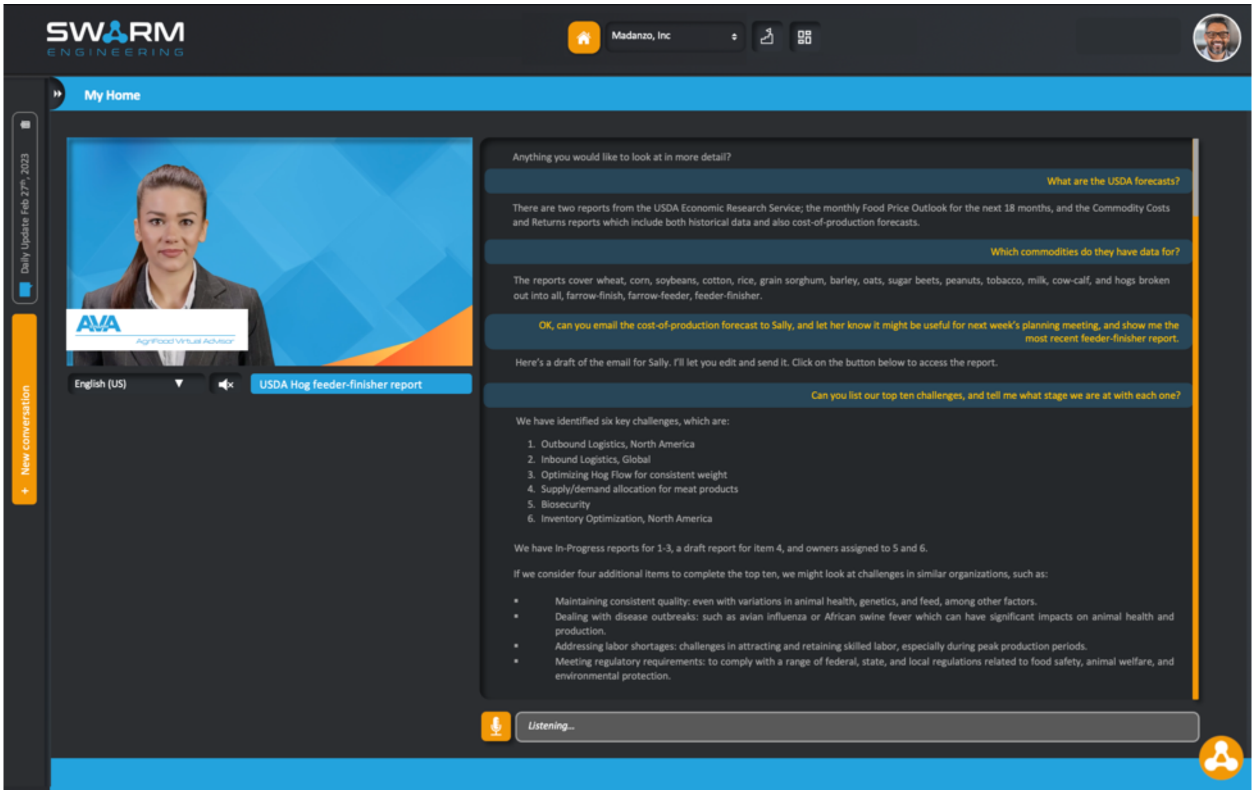

AVA running inside the SWARM platform, helping a C-level exec review challenges

Specialist AI Agents & Managers

For many reasons we believe the next 1-2 years will be challenging for most AgriFood organizations looking to implement LLM systems, although there will also be substantial opportunity for growth. The main causes of difficulty will be:

1. Not taking any action is a high-risk strategy, given the benefits
2. The proliferation of AI & LLM systems will make vendor choices difficult
3. Interaction between AI systems will be very limited
4. The pace of change will be extremely fast
5. Lack of internal (and market) AI expertise will continue

All of the major software vendors have already announced that they have or will soon introduce LLM capabilities into their product ranges. We believe AgriFood specific software vendors will also add LLM and/or digital agent capabilities in the near to mid-term. The challenge is that each will do so independently, as they view their data as one of their primary assets. This will result in many specialized systems or agents with dedicated training and expertise: agent #1 will understand fruit pollination, agent #2 will be a pest expert for blueberries and avocados only, agent #3 will be an agronomist but trained on row crop data, and so on.

A few of the 'mega' vendors in AgriFood, such as John Deere or Monsanto/Bayer may have agents with broader capabilities, but they will likely make recommendations and restrict their advice to products within their closed eco-system, or at the very least show a strong preference for such. That shouldn't pose a problem for organizations bought into that approach, but for AgriFood companies with a more heterogeneous environment, this will prove problematic.

Many (or perhaps most) AgriFood companies themselves will experiment with LLM in some of the primary functions listed earlier, such as customer support, and generating marketing content. They may also build, or attempt to build, customer facing LLM agents of their own.

The Peacock's Tail

A C-level exec interacting with digital agents & employees

The picture that emerges is one we call 'the peacock's tail' in which there will be a rapid proliferation of AI agents within an AgriFood organization. The interactions will be necessary, because each of the agents (especially those from AgriFood vendors) will have specialist, valuable knowledge.

We foresee a scenario where a company may need specific AI agents to manage the AI 'employees'. Interpersonal skills will not be the key job requirement, but an open API and strong interfaces to competing platforms may be necessary. A C-level exec would manage a mix of people, AI agents, and AI managers.

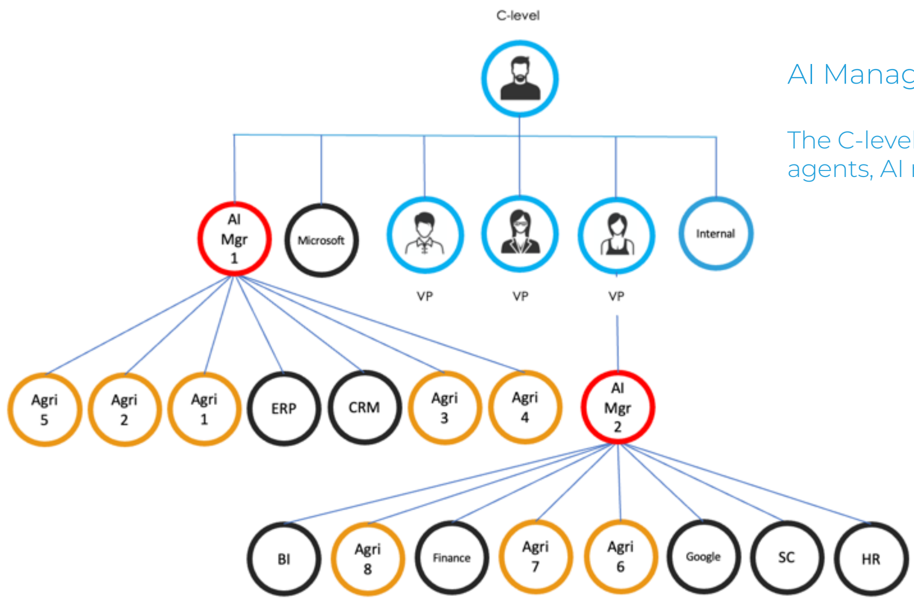

AI Managers arrive

The C-level exec interacts with select agents, AI managers, and employees

There is an attractive alternative, with benefits for many. This would be an **industry association to encourage a common AI approach for LLM**, with a single agent that could work across a range of solutions. Vendors within a market, such as AgriFood, could share the burden of LLM training and develop an agent with expert knowledge across a range of AgriFood subject areas. This would ease the management burden at their target customers, and also deliver a far more powerful solution with broad and deep industry expertise. If you find that thought appealing, feel free to drop me a line: anthony@swarm.engineering

Questions Are King

There is one other critical consideration for AgriFood companies moving into the new era of LLMs, and we are seeing signs of this awareness in new job postings for prompt engineers, or questions analyst. Once you have highly knowledgeable LLM, potentially trained in your industry and your organization, you will need to make questioning a core skill. It is the one thing that AI is not good at doing, and for which humans have a unique capability. Now is the time to consider how good your organization is at asking questions.

Questions govern every aspect of your life. They are the mechanism by which you learn; the tools you use to build relationships. Your questions signal your credibility and knowledge at work and help select your future colleagues during the interview process. The people you encounter in life have varying degrees of training in question theory and practice. Some will know how to extract information from you, whether you wish to surrender that data or not, others will attempt to persuade you to buy goods, join causes, or take actions; all by using questions. You may feel that asking questions is an innate skill; something everybody grasps. You're probably right. This skill may be one of the most profound characteristics that make you human and is something that animals and AI alike cannot do effectively. That doesn't mean you can't improve, though. Most people can swim, but our style, expertise, and technique can vary tremendously. If you want to swim faster, you may have to unlearn some techniques, master new ones, study, train, and even change your diet. Luckily, you don't have to stop eating chocolate to master the art of asking good questions.

Fundamentally, we intuitively know questions are critical in all our communications. We use them at work, at home, and even for internal debates with ourselves. Questions are the primary way that we learn about the world, but equally they can be used to attack or defend, and they reveal far more about our inner thought processes than we might imagine.

In the new world of LLM, your ability to ask the right questions will become even more important. You will have a set of smart agents at your disposal who can answer your many queries, and generate content based on your prompts. **The quality of the LLM responses will depend largely on your expertise at questioning.**

A couple of years ago, I wrote a book called *Questions – A User's Guide*. Strangely enough, it is more relevant today than when it was launched. Here are ten key points that I learned during the years it took me to research and write it:

1. Patterns of questions are so powerful, that they are banned in certain locations (e.g. Prime Minister Question Time in the UK Parliament requires randomized questions)
2. Questions change our memories of events
3. Every industry has a unique questioning methodology; teachers and policeman ask questions in a very different way
4. Modern science is built through formal layers of questions

5. Our biggest regrets in life are frequently the questions we didn't ask
6. We ask the most questions at 4-5 years old – after that, at school and in business, we are rewarded for answering questions, not asking them
7. Everyone is subject to bias; the use of anchoring, order, and acquiescence bias in questions will change the answers you receive
8. AI and trained animals can be very good at answering questions, but they are not capable (yet) of asking, new, creative questions – without human prompts
9. The truth is slippery, and questions won't always help you find it
10. The questions we ask reveal our inner mental models

I would encourage executives to take a look at questioning capability within their company, even before LLMs become more widely available. As with any core competence, training and executive focus & support on this topic can improve this organizational capability.

"When you have AI agents that can answer almost any question about your business, and actually change software systems to deliver dramatically different results, what will you ask them?"

A M Howcroft

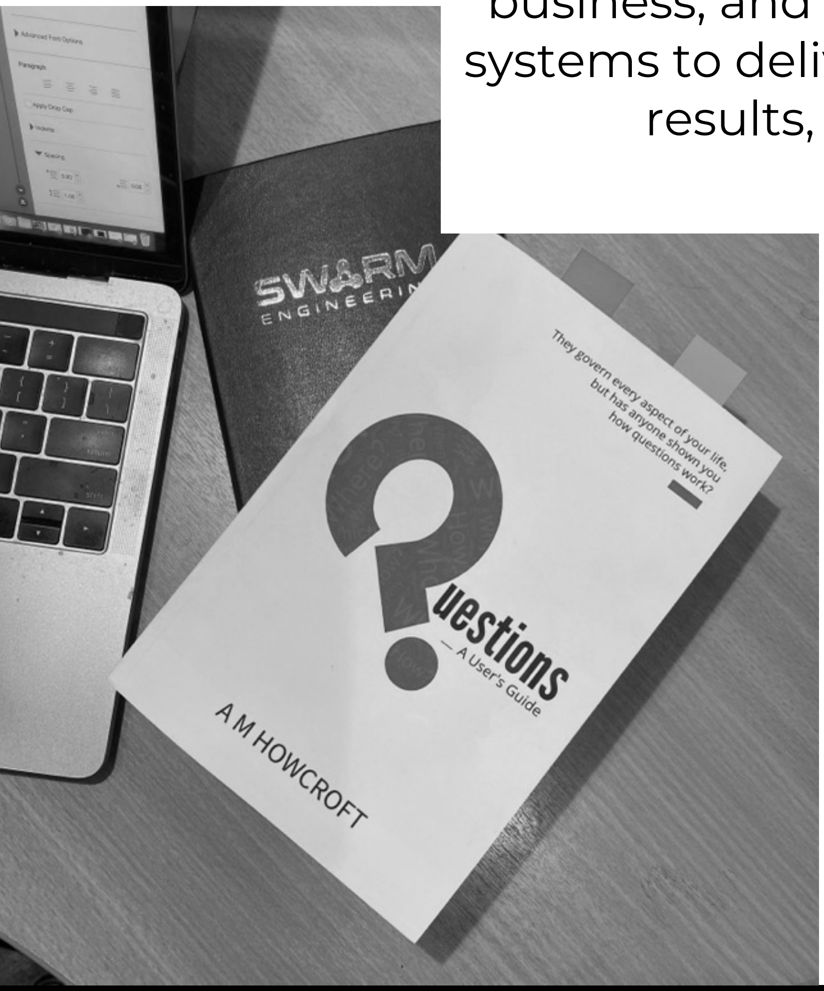

PREDICTIONS

The Backlash

It is likely that there will be a backlash against LLMs in the future, as concerns about privacy, bias, and ethical considerations continue to grow. There will undoubtedly be further reports of unfair or discriminatory outcomes because of the training data and approaches used to generate the LLMs. Additionally, since LLMs require large amounts of data to be effective, there are possible concerns about data privacy and the potential for misuse of personal information, as well as the environmental impact of the computing power needed to create and feed the models. There may also be concerns about the use of LLMs for surveillance.

Finally, as has been widely publicized already, LLMs like other forms of automation, have the potential to replace human workers in certain industries, leading to job losses and economic disruption. Of course, there will undoubtedly be new roles and opportunities, as we have seen during previous major shifts in technology. As with any new technology, it will be important to consider and address any potential risks and concerns as LLMs continue to develop and become more widespread. This will require a collaborative effort between policymakers, industry stakeholders, and the broader public to ensure that LLMs are used in a responsible and ethical manner. By balancing the benefits with the potential risks and concerns, we can ensure that this technology is used to improve our lives and society in a positive way.

For AgriFood companies and vendors alike, the most likely risks in early adoption of LLMs would be a change in consumer behavior, based on negative publicity related to the concerns above. To negate those risks, it will be important for organizations to show they are proactively addressing concerns about privacy, bias, and ethical considerations. There will also be benefits in deploying explainable AI, with levels of transparency over how decisions are made. An additional step organizations can take is to ensure there is human oversight for an AI system. It is one thing to have generative AI create content and write emails, but it would be wise to have a person review these before publication. We should see AI as an advisor and assistant to humans, not a replacement.

Winners and Losers

It is difficult to predict which AgriFood firms will emerge as leaders in the ChatGPT era, as there are many factors that could impact their success. However, firms that are able to rapidly adapt and incorporate LLM and other AI technologies into their operations, collect and manage high-quality data, innovate, and differentiate themselves through the use of AI, collaborate with other organizations, and demonstrate their commitment to sustainability and social responsibility are likely to be more successful than those that do not.

To position themselves for success AgriFood firms should focus on developing and implementing solutions that address key challenges and opportunities in their business, starting with those that have the highest financial or reputational impact. This could involve using LLMs and other AI technologies to improve supply chain optimization, labor management, logistics, animal health monitoring, market intelligence, food safety, soil analysis, sustainability metrics, and so on. By leveraging AI in a responsible and ethical manner and addressing concerns around privacy, bias, and transparency, AgriFood firms can differentiate themselves and gain a competitive edge in an increasingly technology-driven market. This is, of course, a challenging digital transformation while still keeping the existing business running successfully and will require careful navigation.

We believe that making questioning skills a core capability will also bring significant dividends and make working with LLMs more effective. Separating the problem definition from the solution will also be critical, as humans become adept at using questions to define their challenges and rely on the power of AI to solve them. This approach plays to the strengths of both parties and provides a solid foundation for partnership.

The losers may be easier to predict; firms without a solid data collection and quality program will suffer, as will those that decide to sit and wait for too long before at least experimenting with the new capabilities offered by LLMs. Comparisons have been drawn between two other technological shifts – the recent rise (and fall) of crypto currencies, and the introduction of the Internet. While both ChatGPT and crypto rely on complex algorithms and distributed systems to achieve their goals and have the potential to disrupt established industries and business models, crypto was primarily focused on creating decentralized systems for managing and exchanging digital assets. We believe LLMs have more in common with the Internet, as a tool that enables the easier creation and dissemination of information and content on a global scale. Like the Internet, LLMs have the potential to democratize access to information and knowledge, making it possible for people around the world to communicate and collaborate in new ways. Not participating early in the Internet was a decision made by several organizations, such as Blockbuster, Borders, and Sears, and later these decisions would be regarded as a mistake. Companies have a similar choice now.

A New Normal

In two years, the most successful AgriFood companies will likely have adopted LLM and seen several significant benefits:

- ✓ Greater Automation and Efficiency
- ✓ Improved Decision-Making
- ✓ Enhanced Sustainability
- ✓ Improved Customer Experience
- ✓ Greater Customer Loyalty

We believe digital agents such as AVA will be a common feature on the desktops and mobile devices of many employees. These AI agents with LLM skills will have learned about the organization and will be providing increasingly good advice, guidance, and assistance – perhaps even supplanting some of the consultants that were traditionally hired from global System Integrators.

Digital agents will be common on Microsoft Team, Zoom, and Google Meet calls, and may be considerably better than some of their human counterparts at summarizing information and explaining items in a clear and concise manner.

There will be a handful of primary LLM based agents that dominate across technology platforms and/or specific industries, and they will interface with other agents to summarize news and feedback action requests. Some of these dominant agents will be from major software vendors such as Microsoft or Amazon, while others will be formed from industry associations and vendors that have shared learning and technology to create a powerful agent for their market.

New roles and functions will emerge at AgriFood organizations, and they will likely be running challenge centers where questioning experts work with virtual agents to analyze the organization, market, and product opportunities and discover new concepts for efficiency and revenue gains, plus cost savings. Many of these solutions will be assembled and implemented directly from composable applications in a smart market, with no code required. However, the data collection, gathering, and cleaning will still be a bottle neck that holds back new projects. Some things never change.

The first rumblings of 'general AI' may also be taking place, although this is likely still a way into the future, which may be a good thing…

Super Intelligence

Not to finish on a down note, but I read a well-written, unnerving book, called Superintelligence by Nick Hostrum. His thesis is that as machines gain further intelligence, they won't stop at 'Humanville' but will swoosh right past to become far more intelligent. When we think in terms of 'smart' or 'stupid,' our frame of reference is other humans. The author suggests that an AI machine would have significantly greater cognitive capabilities and be "smart in the sense that an average human is smart compared to a beetle or a worm". He points out that a genuinely intelligent computer has several advantages over a human brain: much higher (almost unlimited) storage capacity, and the ability to operate at a speed at least ten thousand times faster than a biological equivalent, allowing it to read a book in a few seconds. If the system were running a million times faster than a biological brain, it could accomplish 1000 years of intellectual work in one working day. It would never forget anything it learned and potentially, could have access to billions of data feeds showing everything from weather to the locations of certain people, health records, stock values, crop data, and much more. An artificial mind of this type would not have the same values as ourselves, unless we programmed them in, and even then, as a learning system it could change its values. Nick looked at how a system like this could emerge. He points out there are multiple ways such a system could be built, such as by a brain emulation or a community of software agents, and that researchers are currently pursuing all avenues. Whichever route is successful, if learning is present, the system will ultimately enter a period of strong recursive growth. Superintelligence could gradually evolve over a period of years or materialize almost instantly—from our perspective. As Nick says, "if and when a take-off occurs, it will likely be explosive". He reviews the various strategies that humans could take to restrict or control this take-off, and then debunks each of them. In simple terms, when a computer system arrives that is super-intelligent, it will out-think or socially engineer us to remove any obstacles we have placed in its way to restrict its growth.

Interestingly, Nick does not predict a terminator-style armageddon where the machines try to eradicate us. As he says, we haven't tried to eliminate ants. More likely, we'd suffer as the system changes the environment to support itself. It might, for example, decide to replace millions of wheat field acres with solar panels, to give itself more energy. There are some fabulous questions in the book about what we want from super-intelligence—how do we avoid catastrophe, and should we hand over the reins to a paternalistic supercomputer? As Nick says, "The point of super-intelligence is not to pander to human preconceptions but to make mincemeat out of our ignorance and folly."

Food for thought, as we move forward with our increasingly powerful LLM systems.

Next Steps – The AgriFood Future Starts Now

To prepare for generative AI and LLM models, AgriFood companies should start by evaluating their current status in the following areas:

- ❑ **Data Maturity**
  Do you have data for all key operational systems? What is the latency? Cleanliness? Accessibility?
- ❑ **AI & Data Science expertise**
  Is there in-house expertise, or do you outsource? Are the team educated on AI capabilities?
- ❑ **Questioning Capability**
  Has there been any formal training on questioning? Is questioning encouraged in the culture?
- ❑ **Critical Process Metrics & Performance**
  Which are the most critical business processes? How much do they cost? Benchmarks vs industry?
- ❑ **Policies for ethics, ESG, data privacy & security**
  LLM will increase the focus on these policies, so check - are they measurable and known?

Beyond this stage, we would recommend the following approach to introducing the new LLM and AI technology into the organization:

- ❑ **Educate**
  Invest in training, and ensure people get hands-on time with LLMs before implementing. Executive sponsorship is critical, so the goals and purpose of the AI is clearly communicated.
- ❑ **Identify & Define Challenges**
  This is a critical phase to catalog potential areas of improvement. Consider creating a Challenge Center, with 2-3 people 'stars' from the business to drive this activity.
- ❑ **Rank Challenges by business impact**
  This will help to prioritize which Challenges to tackle first. You may also want to rank by data maturity, complexity, and so forth to give a balanced view.
- ❑ **Gather missing data**
  No matter how mature your data process, there is likely to be missing data, or information held in heads/spreadsheets that is critical for your high impact challenge. Start the process of capturing this data earlier than you think is required – because it will take longer than expected.
- ❑ **Target a high-impact process, with external assistance**
  Unless you have deep AI & data science expertise, we recommend hiring a vendor that can provide expert assistance to support the roll-out of a high impact challenge solution.
- ❑ **Experiment internally on a non-core process**
  Allow IT and business departments to experiment on non-core processes, to build expertise.
- ❑ **Evaluate internal data for LLM training purposes**
  Review what data could be used to train an internal LLM, to add more value and take expert advice on how to structure this for input. Consider building an internal LLM as an advisor to employees.

Good luck!

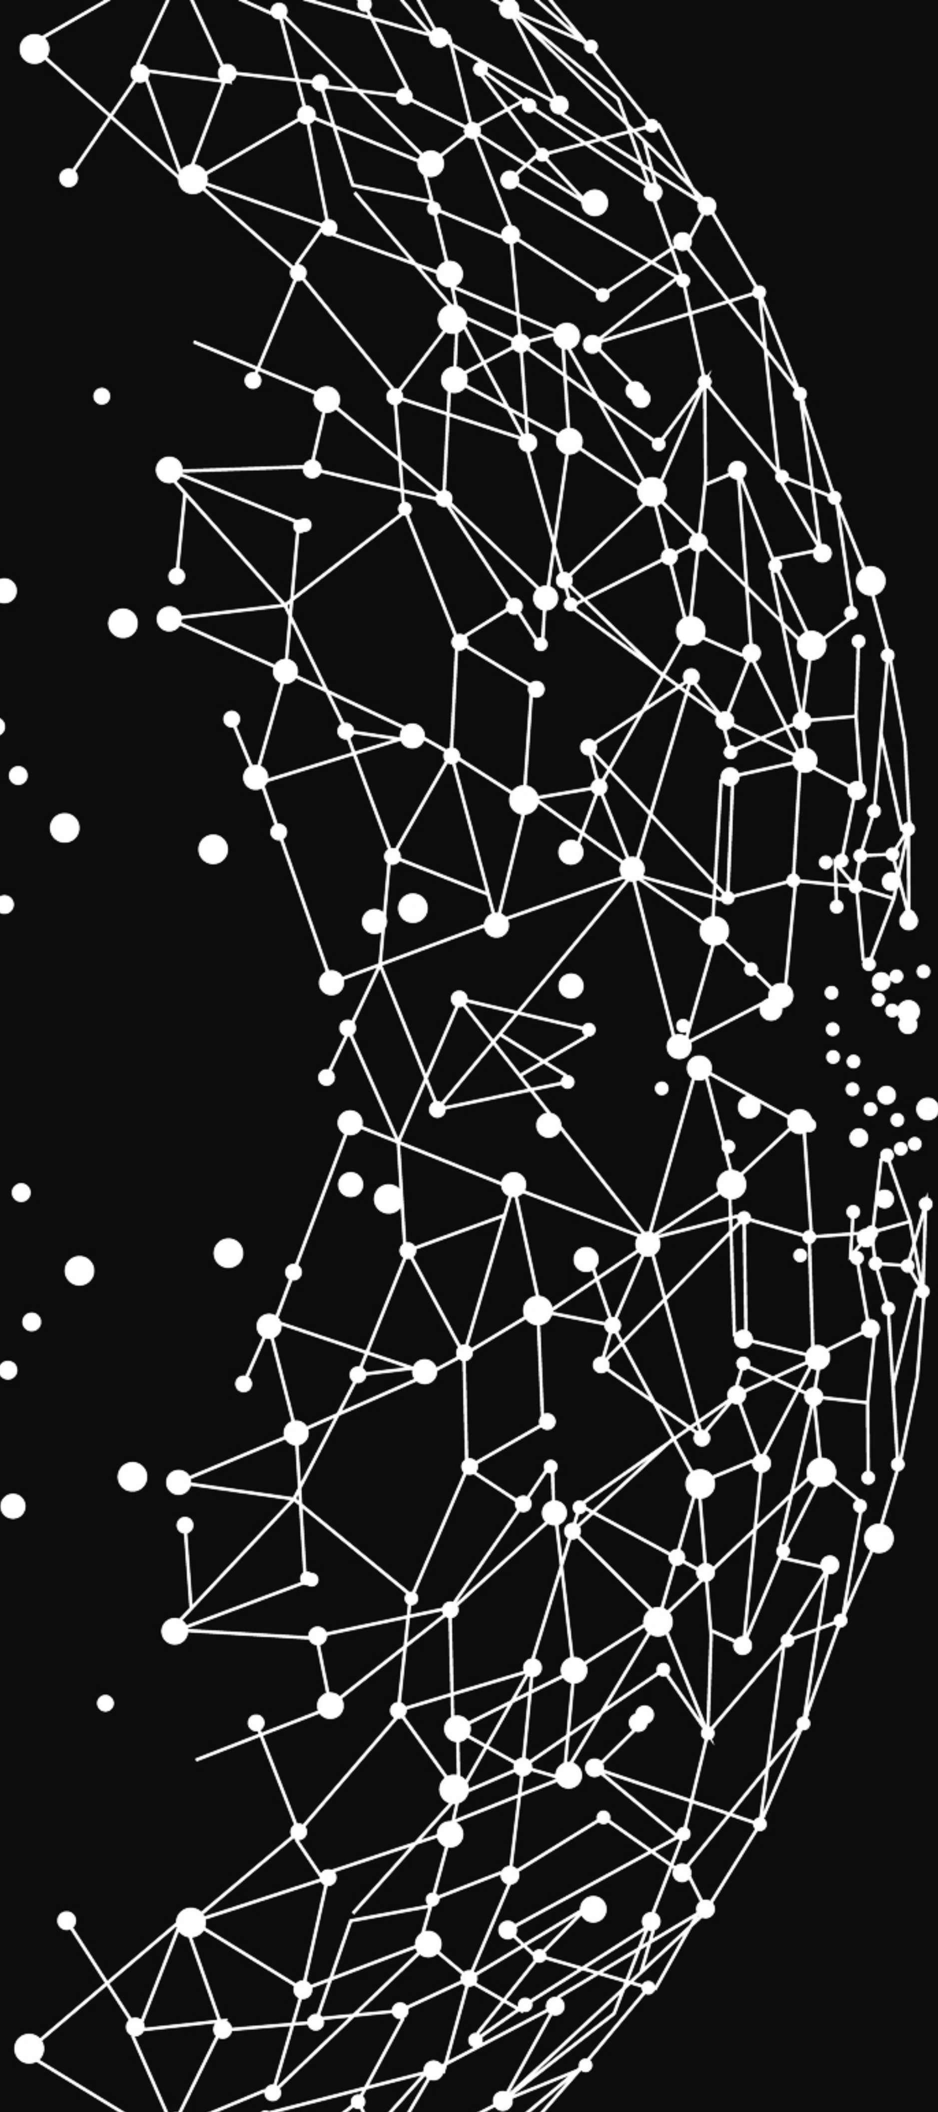

SWARM
ENGINEERING

© SWARM Engineering
March 2023

www.swarm.engineering